selected for the Cb component, then that filter may likely be the optimal filter for the Cr component filter decision as well. We could skip filters that performed poorly for one component's sample when evaluating the other component's sample for the same position. In addition, if the characteristics (e.g. edge or texture classification) of the different color component samples match, then it is more than likely that the same filter would be used for both components. This can considerably reduce complexity and speed up the downsampling process.

[0068] The above process could be performed directly on transfer function domain data (e.g. power law gamma or using a perceptual quantizer-based gamma, e.g. ST.2084/PQ), or directly on linear data.

[0069] The process could involve a 4:4:4 to 4:2:2 downsampling, 4:2:2 to 4:2:0, or 4:4:4 to 4:2:0. For the later case, we can first perform adaptive downsampling across one direction, e.g. horizontally, followed by adaptive downsampling across the other direction, e.g. vertically. It is also possible that the process may involve 2D downsampling directly and the conversion could be done in a single step for both dimensions.

[0070] Filter Optimization

[0071] An encoder may know the filter and transform functions that will be used to upsample the reconstructed image data. Then, with a known upsampling filter, the downsampling filter may be designed to optimize the distortion that will be induced by the upsampling. As previously noted, in an embodiment the encoder may target multiple different delivery systems, e.g. different displays that may have different filters. The known filters for these displays could be considered and evaluated to select the downsampler that results in the best average performance across all the available upsamplers. The various known display upsamplers could be considered with equal weighting, or could be weighted based on the prevalence and importance of each anticipated upsampler.

[0072] Such optimization can be achieved by choosing an optimal filter based on the Cb/Cr component of the original source data, or some other color representation of the original source data (e.g. R, G, and B or X, Y, Z representations). FIG. **6** is a flowchart for an example process for selecting a downsample filter. In the example of FIG. **6**, performing quantization and downsampling on the Cb/Cr component and then performing the inverse functions **620** to reconstruct the Cbr/Crr component using each different available filter, the original Cb/Cr may be compared to the reconstructed Cbr/Crr data to calculate the distortion added during the processing. The data can optionally be converted back to R, G, and B, or X, Y, Z, and the distortion be computed in those spaces instead. The filter that best minimizes distortion may then be selected for use in the encoding process.

[0073] As shown in FIG. **6**, initially for a portion of the image data, the Cb and Cr components of the portion are encoded using one of the available upsampling filters (block **610**). Then the reconstructed Crr and Cbr components of the portion are created using the corresponding downsampling filter (block **620**). Then the distortion introduced in the encoding and reconstruction process may be calculated by comparing the original Cr/Cb components to the reconstructed Crr/Cbr components (block **630**). Once all the available filters have been considered, the filter that results in the lowest distortion will be selected (block **640**). The distortion metric used in the encoding decision process may be, for example, the mean or sum of absolute differences (MAD or SAD), the sum of absolute transformed differences (SATD), the mean or sum of square differences/errors (MSE or SSE), the peak signal to noise ratio (PSNR), the structural similarity index (SSIM), or other suitable operations.

[0074] Filter optimization may also occur end to end in different domains. For example, after Cbr and Crr are calculated, data in the RGB or XYZ color space may be calculated based on each of the original YCbCr and the reconstructed YrCbrCrr. Then the filter that minimizes distortion XYZ and XYZr may be calculated. This approach requires that we first make a preliminary decision for one color component, e.g. Cr, and using that decision try to optimize the other color component, i.e. Cb. After the optimization of this color component is completed we can go back to Cr and refine it further given the initial joint decision for Cb. This process may iterate multiple times. Iteration may repeat until no further change in performance is observed, or after N steps that may be determined by the user or application. The application, for example, may determine the number of steps adaptively, based, for example, on overall complexity budget involving the downsampling/upsampling process. The Y component would also need to be considered in this process, whereas characteristics and quantization on the Y component, as discussed in other sections of this application, may also be considered. Distortion in this extended color space could again be based on MAD, SAD, SATD, SSE, or PSNR. Distortion could be computed in the linear space or using an appropriately defined transfer function. Distortion from each color component in this extended space could be combined using an average or a weighted average, based on an importance of each color component, to generate a single distortion measurement. Weighting could also include content characteristics such as local and global brightness, whether the sample corresponds or is near an edge, texture, or a flat region among others.

[0075] Filters may be chosen based on comparisons of an original and a reconstructed sequence, scene, image, slice, block, pixel, or other portion of image data.

[0076] Similarly, with known filters, the color transform matrix (M) can be optimized in a similar manner as described above to minimize distortion.

[0077] With known upsampling and downsampling filters, and when a fixed bit-depth for quantization is used, the quantization parameters used in the quantization transform Q( ) may be optimized in a similar manner as described to minimize distortion. Such optimization may also take into consideration quantization boundaries and deadzones.

[0078] Information about the selected optimizations may be sent from the encoder to a receiver or decoder as part of encoding hints data transmitted with the encoded data. Encoding hints may be transmitted as side information and included in a bitstream such as H.264 as a Supplemental Enhancement Information (SEI) message. SEI messages may include a Post Filter Hint SEI message or a Chroma resampling filter hint SEI message. Such encoding hints data may include the transform matrix used in color transform, the parameters used in the quantization transform, and information about the filters used in the optimization process among others. A more complex process may also include aspects of the actual video encoding and decoding, i.e., through an initial coding pass using a codec such as H.264